

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims

1 1. (currently amended) A mass storage controller system, comprising:
2 a plurality of controllers for controlling an array of storage devices, each of the plurality
3 of controllers comprising:
4 a CPU for controlling the operation of a controller;
5 program memory, coupled to the CPU, for storing program instructions and variables for
6 the operation of the CPU; and
7 cache memory, coupled to the CPU, for storing information related to the array of
8 storage devices;
9 wherein a controller of the plurality of controllers initiates a task to be performed, the
10 controller initiating the task establishes a task coordination data object shared by the plurality of
11 controllers, wherein the task coordination data object ~~consists of~~ comprises discrete partitions of
12 the task comprising task instructions and states for each partition of the task, task instructions for
13 each partition of the task capable of being completed separately by one of the plurality of
14 controllers ~~to allow the task to be completed by way of the cumulative effort of the plurality of~~
15 ~~controllers completing separately the partitions of the task~~, and wherein a free controller of the
16 plurality of controllers selects a partition of the task available for completing separately and
17 independently of the other controllers as indicated by the states for each partition of task
18 instructions.

1 2. (Original) The mass storage controller system of claim 1, wherein the state indicates
2 whether a partition is READY, IN PROGRESS, or COMPLETE.

1 3. (Original) The mass storage controller system of claim 2, wherein a controller selects
2 a partition by examining the partitions in a READY state and selecting at least one partition in
3 the READY state to operate on.

1 4. (Original) The mass storage controller system of claim 3, wherein a partition is in an
2 IN PROGRESS state during processing.

1 5. (Original) The mass storage controller system of claim 4, wherein a controller sets the
2 partition selected for processing to a COMPLETE state upon completion of processing for a
3 partition.

1 6. (Original) The mass storage controller system of claim 1, wherein a controller selects
2 a partition by examining the partitions in a READY state and selecting at least one partition in
3 the READY state to operate on.

1 7. (Original) The mass storage controller system of claim 1, wherein a partition is in an
2 IN PROGRESS state during processing.

1 8. (Original) The mass storage controller system of claim 1, wherein a controller sets the
2 partition selected for processing to a COMPLETE state upon completion of processing for a
3 partition.

1 9. (Original) The mass storage controller system of claim 1, wherein the states provide a
2 semaphore-mechanism for allowing a controller to ascertain whether to acquire control over a
3 partition.

1 10. (Original) The mass storage controller system of claim 1, wherein the initiating
2 controller is notified when all partition states are COMPLETE and performs whatever
3 completion actions are required.

1 11. (Original) The mass storage controller system of claim 1, wherein the task
2 coordination data object includes information about an operation to be performed and a data set
3 to be operated on.

1 12. (currently amended) A mass storage array subsystem, comprising:
2 a plurality of storage devices;
3 a backplane, coupled to the plurality of storage devices, adapted to couple to said
4 plurality of storage devices; and
5 a plurality of controllers, coupled to the backplane, for controlling the plurality of
6 storage devices, the plurality of controllers having a first interface coupled to a host system and a
7 second interface coupled to said backplane to communicate with said plurality of storage
8 devices;
9 wherein each of the plurality of controllers comprise a CPU for controlling the operation
10 of a controller, program memory for storing program instructions and variables for the operation
11 of the CPU and cache memory for storing information related to the array of storage devices, and
12 wherein a controller of the plurality of controllers initiates a task to be performed, the controller
13 initiating the task establishes a task coordination data object shared by the plurality of
14 controllers, wherein the task coordination data object ~~consists of~~ comprises discrete partitions of
15 the task comprising task instructions and states for each partition of the task, task instructions for
16 each partition of the task capable of being completed separately by one of the plurality of
17 controllers ~~to allow the task to be completed by way of the cumulative effort of the plurality of~~
18 ~~controllers completing separately the partitions of the task~~, and wherein a free controller of the
19 plurality of controllers selects a partition of the task available for completing separately and
20 independently of the other controllers as indicated by the states for each partition of task
21 instructions.

1 13. (Original) The mass storage array subsystem of claim 12, wherein the state indicates
2 whether a partition is READY, IN PROGRESS, or COMPLETE.

1 14. (Original) The mass storage array subsystem of claim 13, wherein a controller selects
2 a partition by examining the partitions in a READY state and selecting at least one partition in
3 the READY state to operate on.

1 15. (Original) The mass storage array subsystem of claim 14, wherein a partition is in an
2 IN PROGRESS state during processing.

1 16. (Original) The mass storage array subsystem of claim 15, wherein a controller sets
2 the partition selected for processing to a COMPLETE state upon completion of processing for a
3 partition.

1 17. (Original) The mass storage array subsystem of claim 12, wherein a controller selects
2 a partition by examining the partitions in a READY state and selecting at least one partition in
3 the READY state to operate on.

1 18. (Original) The mass storage array subsystem of claim 12, wherein a partition is in an
2 IN PROGRESS state during processing.

1 19. (Original) The mass storage array subsystem of claim 12, wherein a controller sets
2 the partition selected for processing to a COMPLETE state upon completion of processing for a
3 partition.

1 20. (Original) The mass storage array subsystem of claim 12, wherein the states provide
2 a semaphore-mechanism for allowing a controller to ascertain whether to acquire control over a
3 partition.

1 21. (Original) The mass storage array subsystem of claim 12, wherein the initiating
2 controller is notified when all partition states are COMPLETE and performs whatever
3 completion actions are required.

1 22. (Original) The mass storage array subsystem of claim 12, wherein the task
2 coordination data object includes information about an operation to be performed and a data set
3 to be operated on.

1 23. (currently amended) A method for cooperative distributed task management in a
2 storage subsystem with multiple controllers, comprising:
3 initiating by an initiating controller a task to be performed;
4 establishing by the initiating controller a task coordination data object shared by the
5 multiple controllers, wherein the task coordination data object ~~consists of~~ comprises discrete
6 partitions of the task comprising task instructions and states for each partition of the task, task
7 instructions for each partition of the task capable of being completed separately by one of the
8 plurality of controllers ~~to allow the task to be completed by way of the cumulative effort of the~~
9 ~~plurality of controllers completing separately the partitions of the task; and~~
10 selecting by a free controller a partition of the task available for completing separately
11 and independently of the other controllers as indicated by the states for each partition of task
12 instructions.

1 24. (Original) The method of claim 23 further comprising indicating a state of a partition
2 as being READY, IN PROGRESS, or COMPLETE.

1 25. (Original) The method of claim 24 wherein the selecting by a free controller is
2 performed by examining the partitions in a READY state and selecting at least one partition in
3 the READY state to operate on.

1 26. (Original) The method of claim 25, wherein a partition is in an IN PROGRESS state
2 during processing.

1 27. (Original) The method of claim 26 further comprising setting by a controller a
2 partition selected for processing to a COMPLETE state upon completion of processing for the
3 partition.

1 28. (Original) The method of claim 23, wherein the selecting by a free controller is
2 performed by examining the partitions in a READY state and selecting at least one partition in
3 the READY state to operate on.

1 29. (Original) The method of claim 23, wherein a partition is in an IN PROGRESS state
2 during processing.

1 30. (Original) The method of claim 23 further comprising setting by a controller a
2 partition selected for processing to a COMPLETE state upon completion of processing for the
3 partition.

1 31. (Original) The method of claim 23 wherein the states provide a semaphore-
2 mechanism for allowing a controller to ascertain whether to acquire control over a partition.

1 32. (Original) The method of claim 23 further comprising notifying the initiating
2 controller when all partition states are complete and performing completion actions that are
3 required.

1 33. (Original) The method of claim 23, wherein the task coordination data object
2 includes information about an operation to be performed and a data set to be operated on.

1 34. (currently amended) An article of manufacture comprising a program storage
2 medium readable by a computer, the medium tangibly embodying one or more programs of
3 instructions executable by the computer to perform a method for cooperative distributed task
4 management in a storage subsystem with multiple controllers, the method comprising:
5 initiating by an initiating controller a task to be performed;
6 establishing by the initiating controller a task coordination data object shared by the
7 multiple controllers, wherein the task coordination data object ~~consists of~~ comprises discrete
8 partitions of the task comprising task instructions and states for each partition of the task, task
9 instructions for each partition of the task capable of being completed separately by one of the
10 plurality of controllers ~~to allow the task to be completed by way of the cumulative effort of the~~
11 ~~plurality of controllers completing separately the partitions of the task; and~~
12 selecting by a free controller a partition of the task available for completing separately
13 and independently of the other controllers as indicated by the states for each partition of task
14 instructions.

1 35. (Original) The article of manufacture of claim 34 further comprising indicating a
2 state of a partition as being READY, IN PROGRESS, or COMPLETE.

1 36. (Original) The article of manufacture of claim 35 wherein the selecting by a free
2 controller is performed by examining the partitions in a READY state and selecting at least one
3 partition in the READY state to operate on.

1 37. (Original) The article of manufacture of claim 36, wherein a partition is in an IN
2 PROGRESS state during processing.

1 38. (Original) The article of manufacture of claim 37 further comprising setting by a
2 controller a partition selected for processing to a COMPLETE state upon completion of
3 processing for the partition.

1 39. (Original) The article of manufacture of claim 34, wherein a free controller selects a
2 partition by examining the partitions in a READY state and selecting at least one to operate on.

1 40. (Original) The article of manufacture of claim 34 further comprising notifying the
2 initiating controller is notified when all partition states are complete and performs whatever
3 completion actions required.

1 41. (Original) The article of manufacture of claim 34, wherein the task coordination data
2 object includes information about an operation to be performed and a data set to be operated on.

1 42. (currently amended) A mass storage controller system, comprising:
2 a plurality of intermediate controller means for controlling an array of storage devices,
3 each of the plurality of controllers comprising:
4 CPU means for controlling the operation of a controller;
5 program memory means, coupled to the CPU means, for storing program instructions
6 and variables for the operation of the CPU; and
7 cache memory means, coupled to the CPU means, for storing information related to the
8 array of storage devices;
9 wherein an intermediate controller means of the plurality of intermediate controller
10 means initiates a task to be performed, the intermediate controller means initiating the task
11 establishes a task coordination data object shared by the plurality of intermediate controller
12 means controller means, wherein an intermediate controller means of the plurality of
13 intermediate controller means initiates a task to be performed, the intermediate controller means
14 initiating the task establishes a task coordination data object shared by the plurality of
15 intermediate controller means, wherein the task coordination data object ~~consists of~~ comprises
16 discrete partitions of the task comprising task instructions and states for each partition of the
17 task, task instructions for each partition of the task capable of being completed separately by one
18 of the plurality of intermediate controller means ~~to allow the task to be completed by way of the~~
19 ~~cummulative effort of the plurality of intermediate controller means completing seapratly the~~
20 ~~partitions of the task~~ and wherein a free intermediate controller means selects a partition of the
21 task available for completing separately and independently of the other controllers as indicated
22 by the states for each partition of task instructions.